arXiv:1610.07499v1 [cs.CC] 24 Oct 2016

Dynamic Complexity of the Dyck Reachability

Patricia Bouyer and Vincent Jugé
LSV, CNRS & ENS Cachan, Univ. Paris-Saclay, France
This work is supported by EU under ERC EQualIS (FP7-308087).

October 25, 2016

Abstract

Dynamic complexity is concerned with updating the output of a problem when the input is slightly changed. We study the dynamic complexity of Dyck reachability problems in directed and undirected graphs, where updates may add or delete edges. We show a strong dichotomy between such problems, based on the size of the Dyck alphabet. Some of them are P-complete (under a strong notion of reduction) while the others lie either in DynFO or in NL.

1 Introduction

Dynamic problems and dynamic complexity. In this paper, we focus on the dynamic complexity of some reachability problems. Standard complexity theory aims at developing algorithms that, given an input of some problem, compute an output as efficiently as possible. Its dynamic variant is focused on algorithms that are capable of efficiently updating the output after a small change of the input [13, 10, 14]. Such algorithms may rely on auxiliary data about the current instance of the problem, and update it when the instance is modified.

A well-studied dynamic complexity class is DynFO. An algorithm is in DynFO if the output and the auxiliary data can be updated by FO formulas after a small change of the input. Variants of DynFO include the class DynFO^+, which allows polynomial-time precomputations, and DynTC^0, in which updates of the auxiliary data are performed by TC^0 circuits.

Consider the problem of reachability in directed graphs, and update operations that consist in inserting or deleting edges (one at a time). It was recently proven that this problem belongs to the class DynFO [2], which had been conjectured for decades.

Furthermore, like static complexity classes, dynamic complexity classes come with natural notions of reduction. The class DynFO is closed under bounded expansion first-order reductions (hereafter called bfo reductions), which are specific L reductions (L is for logarithmic space). A bfo reduction from a problem to another one is a first-order mapping from instances of the first problem to instances of the latter one, such that performing an update operation on the instance of the first problem amounts to performing a bounded number of update operations on the instance of the latter problem. Similarly, the class DynFO^+ is closed under bounded expansion first-order reductions with polynomial-time precomputation (hereafter called bfo$^+$ reductions).

Reachability problems and language theory. Dyck reachability problems lie at the interface between two areas. On the one hand, language theory is concerned with handling descriptions of languages, that is sets of words, with respect to various questions: Is a language empty, finite or infinite? What about the intersection or the union of two language? Does a language contain a given word? Among the best known and most simple classes of languages are regular and context-free languages. On the other hand, reachability problems deal with the existence of paths in graphs, and include questions such as: Does there exist a path between two given vertices? How long must be such paths?

Dyck reachability problems are focused on the existence of paths in labeled graphs, whose labels belong to a given *Dyck* language. Dyck languages are languages of well-parenthesized words and, roughly speaking, are the most simple context-free languages that are not regular. The Dyck

reachability problem in labeled directed acyclic graphs was proven to be in DynFO [14], when considering two types of update operations on labeled graphs, which are insertion and deletion of edges. Whether this result extends to all labeled directed graphs was then an open question.

Our contributions. We study this open question, and we distinguish the Dyck reachability problem in two different ways. Is the labeled graph directed or undirected? How many symbols does the Dyck alphabet contain?

We prove that there exists a strong dichotomy between the dynamic complexity of such problems, based on the size of the Dyck alphabet. In the case of a unary Dyck alphabet, the Dyck reachability problem lies in NL (non-deterministic logarithmic space), and even lies in DynFO in the case of undirected graphs; this contrasts with the case of binary Dyck alphabets, where we prove that the Dyck reachability problem is P-complete under bfo^+ reductions. Furthermore, it is widely believed [13] that no P-complete problems under bfo^+ reductions lie in classes such as DynFO or the slightly broader class DynFO^+.

Related works. From its very inception 20 years ago, dynamic complexity has been a framework of study for serveral variants of reachability problems and language theory problems. The class DynFO was shown to contain reachability problems in directed acyclic graphs [4], undirected graphs [13] and, most recently, in all directed graphs [2]; regular and Dyck languages [13], then all context-free languages [5]; Dyck reachability in directed acyclic graphs [14].

At the same time, finding natural problems that are NL- or P-complete (under L reductions) and belong to low dynamic complexity classes such as DynFO, DynFO^+ or DynTC^0 is an ongoing challenge. All known P-complete problems lying in DynFO rely on highly redundant inputs, hence may be seen as artificial [13]. Hence, a notion of *non-redundant projection* [12] was introduced. Non-redundant projections are a special kind of P reductions, which contains, in particular, bfo and bfo^+ reductions.

Hence, for every static complexity class $\mathcal{C}$, we define *non-redundant* $\mathcal{C}$-complete problems as those problems that are $\mathcal{C}$-complete both under L reductions and under non-redundant projections. Most canonical P-complete problems are non-redundant, hence non-redundancy may be seen as a prerequisite for being a "natural" problem.

A breakthrough was the proof that the Dyck reachability problem in acyclic directed graphs, which is a non-redundant LogCFL-complete problem, belongs to DynFO [14]. Yet, Weber and Schwentick also observed in [14] that "it remains open whether there is a non-redundant P-complete problem that allows efficient updates. The Dyck reachability problem might be a candidate." In this paper, we prove that this latter claim is unlikely to hold, even if we restrict ourselves to the seemingly easier case of undirected graphs.

2 Definitions

2.1 Dyck reachability problems

A labeled directed graph is a triple $G = (V, L, E)$ where V is a finite set of vertices, L is a finite set of labels and $E \subseteq V \times L \times V$ is a finite set of edges. The graph G is said to be *unlabeled* if L is a singleton set; in that case, we may directly represent G as a pair (V, E) where $E \subseteq V \times V$. The graph G is also said to be *undirected* if the relation E is symmetric, i.e. if, for every edge (v, θ, w) in E, the triple (w, θ, v) also belongs to E.

A path in the graph G is a finite sequence of edges $\pi = (v_1, \theta_1, w_1) \cdot (v_2, \theta_2, w_2) \cdot \ldots \cdot (v_k, \theta_k, w_k)$ such that $v_{i+1} = w_i$ for all $i \in \{1, \ldots, k-1\}$. The vertex v_1 is called the *source* of π, and w_k is called the *sink* of π. We also denote by $\lambda(\pi)$ the word $\theta_1 \cdot \ldots \cdot \theta_k$, which is called the *label* of π.

Assume that the label set L is of the form $L = \{\ell_1, \ldots, \ell_n\} \uplus \{\overline{\ell}_1, \ldots, \overline{\ell}_n\}$ for some integer $n \geq 1$. The *Dyck language* associated with L is the context-free language $\mathbf{D}_n$ built over the grammar: $S \to \varepsilon \mid \ell_1 \cdot S \cdot \overline{\ell}_1 \cdot S \mid \ldots \mid \ell_n \cdot S \cdot \overline{\ell}_n \cdot S$, where ε is the empty word. The set $\{\ell_1, \ldots, \ell_n\}$ is said to be the *Dyck alphabet* of that language.

The *n-letter Dyck reachability problem* asks whether, given two vertices s and t of G, there exists a path in G, with source s and sink t, and whose label belongs to the Dyck language $\mathbf{D}_n$

(the actual value of the label set L does not matter, as long as its elements can be partitioned in n ordered pairs). The n-letter *undirected* Dyck reachability problem is the restriction of that problem to the case where the underlying graph G is constrained to be undirected.

2.2 Dynamic complexity

In this paper, we study the dynamic complexity of Dyck reachability problems. To that end, we first introduce briefly the formalisms of descriptive and dynamic complexity here, and refer to [13, 11, 9] for more details.

Descriptive complexity aims at characterizing positive instances of a problem using logical formulas. The input is then described as a logical structure described by a set of k-ary predicates (the *vocabulary*) over its universe. For example, a graph can be described by a ternary predicate representing its edges, with the set of states and of labels (usually identified with $\{1, \dots, n\}$ for some n) as the universe. The problem of deciding whether some state has at most one outgoing edge can be described by the first-order formula $\exists x.\forall y.\forall z.(E(x,y) \wedge E(x,z)) \Rightarrow (y = z)$. The class FO contains all problems that can be characterized by such first-order formulas. This class corresponds to the circuit-complexity class AC^0 (under adequate reductions) [1].

Dynamic complexity aims at developing algorithms that can efficiently update the output of a problem when the input is slightly changed, for example reachability of one vertex from another one in a graph. We would like our algorithm to take advantage of previous computations in order to very quickly decide the existence of a path in the modified graph.

Formally, a decision problem S is a subset of the set of τ-structures $\mathsf{Struct}(\tau)$ built on a vocabulary τ. In order to turn S into a dynamic problem DynS, we need to define a finite set of allowed updates. For instance, we might use a 2-ary operator $\mathsf{ins}(x, y)$ that would insert an edge between nodes x and y. For a universe of size n, the set of update operations forms a finite alphabet, denoted by Σ_n. A finite word in Σ_n^* then corresponds to a structure obtained by applying a sequence of update operations of Σ_n to the empty structure $\mathcal{I}_n$ over the vocabulary τ. The language DynS_n is defined as the set of those words in Σ_n^* that correspond to structures of S, and DynS is the union (over all n) of all such languages.

A dynamic machine is a uniform family $(M_n)_{n\in\mathbb{N}}$ of deterministic finite automata $M_n = \langle Q_n, \Sigma_n, \delta_n, s_n, F_n\rangle$ over an update alphabet Σ_n, with an update transition function δ_n. The set of states can be encoded as a structure over some vocabulary τ^{aux}, which contains the vocabulary τ, and corresponds to a polynomial-size auxiliary data structure. Such a machine solves a dynamic problem if $\mathsf{DynS}_n = \mathcal{L}(M_n)$ for all n. It is in the dynamic complexity class $\mathcal{C}'$-$\mathsf{Dyn}\mathcal{C}$ (or simply $\mathsf{Dyn}\mathcal{C}$ if $\mathcal{C} = \mathcal{C}'$) if the update transition function and accepting set can be computed in $\mathcal{C}$, while the initial state can be computed in $\mathcal{C}'$. In other words, solving the initial instance of the problem can be done in $\mathcal{C}'$, and after any update of the input (specified by some letter of Σ_n), further calculations to solve the problem on that new instance are restricted to the class $\mathcal{C}$. Of course, for the a dynamic complexity class $\mathcal{C}'$-$\mathsf{Dyn}\mathcal{C}$ to have some interest, the class $\mathcal{C}$ should be easier than the static complexity class of the original problem.

In this paper, we only consider the case where $\mathcal{C} = \mathsf{FO}$, and where $\mathcal{C}' = \mathsf{FO}$ or $\mathcal{C}' = \mathsf{P}$, meaning that first-order formulas will be used to describe how predicates are updated along transitions, and that we may make use of polynomial-time precomputations. As a convention, we will denote the class P-DynFO by DynFO^+, and we recall that the simple notation DynFO is for FO-DynFO.

2.3 Dynamic reductions

Dynamic complexity comes with the notion of dynamic reductions [13]. Let $\mathcal{C}$ be a complexity class. A (static) $\mathcal{C}$ reduction from a decision problem $\mathcal{P}$ to another decision problem $\mathcal{Q}$ is a mapping in $\mathcal{C}$ from the instances of $\mathcal{P}$ to the instances of $\mathcal{Q}$ that associates every positive instance of $\mathcal{P}$ with a positive instance of $\mathcal{Q}$, and every negative instance of $\mathcal{P}$ with a negative instance of $\mathcal{Q}$. Standard P-completeness results use L reductions [6].

A *dynamic reduction* from a dynamic problem $\mathcal{P}$ (with vocabulary τ_1) to another dynamic problem $\mathcal{Q}$ (with vocabulary τ_2) is a mapping from $\mathsf{Struct}(\tau_1)$ to $\mathsf{Struct}(\tau_2)$ such that:

- every positive (respectively, negative) instance of $\mathcal{P}$ is mapped to a positive (respectively, negative) instance of $\mathcal{Q}$;

- every update on an instance i_1 of $\mathcal{P}$ results in a *well-behaved* sequence of updates on the instance i_2 of $\mathcal{Q}$ to which i_1 is mapped.

Dynamic reductions have therefore several parameters: the complexity class to which the mapping belongs, and the sequences of updates that are allowed.

The dynamic classes DynFO and DynFO$^+$ are respectively closed under bounded expansion first-order (bfo for short) and bounded expansion first-order with polynomial-time precomputation (bfo$^+$ for short) reductions [13]. A dynamic reduction μ from $\mathcal{P}$ to $\mathcal{Q}$ is bfo$^+$ if it is a FO reduction and if every update on an instance i_1 of $\mathcal{P}$ results in a bounded sequence of FO updates on its image $\mu(i_1)$. If, furthermore, the empty structure $\mathcal{I}_1$ is mapped to a structure $\mu(\mathcal{I}_1)$ that can be obtained by applying a bounded sequence of FO updates on the empty structure $\mathcal{I}_2$, then we say that μ is bfo.

Note that dynamic reductions can be applied to the class P (which coincides with the class DynP, under the assumption that updates are one-bit input changes). So, being P-hard for bfo$^+$ reductions is a priori stronger than being P-hard for L reductions. Furthermore, it is known that the classes of bfo and of bfo$^+$ reductions are closed under composition and that the circuit value problem is a P-complete problem for bfo$^+$ reductions [13]. Hence, every P problem to which the circuit value problem is bfo$^+$-reducible is also P-complete problem for bfo$^+$ reductions.

2.4 Main result

We are now in a position to formally present our main result.

Theorem 1. *The 1-letter Dyck reachability problem is in* NL*, and the 1-letter undirected Dyck reachability problem is in* NL $\cap$ Dyn-FO*. Furthermore, for all integers $n \geq 2$, both the n-letter Dyck reachability problem and the n-letter undirected Dyck reachability problems are* P*-complete for bfo$^+$ reductions.*

Remark 2. Note that NL $\cap$ Dyn-FO is not known to be strictly included in NL. Nevertheless, the case of undirected graph appears to be "easier" than the case of directed graphs in the 1-letter case. Hence, the P-hardness of both cases for alphabets with at least two letters appears rather unexpected.

3 One-letter (undirected) Dyck reachability problems

We prove here the first part of Theorem 1, that is we assume $n = 1$. We first observe that the 1-letter Dyck reachability problem is equivalent to a standard reachability problem in one-counter automata (without zero-tests), which is known to belong to NL [3, 7]. The 1-letter undirected Dyck reachability problem is a restriction of the 1-letter undirected Dyck reachability problem, hence it is in NL as well. Furthermore, we make the following claim.

Proposition 3. *Let s and t be two distinct vertices of an undirected labeled graph $G = (V, E, L)$, with $L = \{\ell_1, \overline{\ell}_1\}$. There exists a Dyck path from s to t in G if and only if:*

- *the set $\{x \in V \mid (s, \ell_1, x) \in E\}$ is non-empty;*
- *the set $\{y \in V \mid (t, \overline{\ell}_1, y) \in E\}$ is non-empty;*
- *there exists a path of even length from s to t in G.*

Proof. First, if there exists a Dyck path $\pi = (v_i, \lambda_i, v_{i+1})_{0 \leq i < k}$ with $s = v_0$ and $t = v_k$, then $\lambda_0 = \ell_1$, $\lambda_{k-1} = \overline{\ell}_1$, and the sets $\{0 \leq i < k \mid \lambda_i = \ell_1\}$ and $\{0 \leq i < k \mid \lambda_i = \overline{\ell}_1\}$ have the same cardinality, which proves that k is an even number.

Conversely, assume that $s \neq t$ and that the three conditions of Proposition 3 hold. Let $\pi = (v_i, \lambda_i, v_{i+1})_{0 \leq i < 2k}$ be a path of length $2k$ from s to t in G, for some integer $k \geq 1$. Let κ be the cardinality of the set $\{0 \leq i < 2k \mid \lambda_i = \ell_1\}$ and let $\overline{\kappa}$ be the cardinality of the set $\{0 \leq i < 2k \mid \lambda_i = \overline{\ell}_1\}$. Since $\kappa + \overline{\kappa} = 2k$, we have $\overline{\kappa} - \kappa = 2(k - \kappa)$.

Furthermore, consider vertices $x, y \in V$ such that (s, ℓ_1, x) and $(t, \overline{\ell}_1, y)$ belong to E. Since the graph is undirected, there exist also edges (x, ℓ_1, s) and $(y, \overline{\ell}_1, t)$. Let ρ_1 be the length-2 circuit

$(s, \ell_1, x) \cdot (x, \ell_1, s)$, and let ρ_2 be the length-2 circuit $(t, \overline{\ell}_1, y) \cdot (y, \overline{\ell}_1, x)$. One checks easily that the path $\rho_1^k \cdot \pi \cdot \rho_2^\kappa$ is a Dyck path in G, where ρ_1^k is the concatenation of k occurrences of ρ_1, and ρ_2^κ is the concatenation of κ occurrences of ρ_2. □

Hence, checking whether there exists a Dyck path from s to t in G amounts to checking whether $s = t$ or, if $s \neq t$, whether the sets $\{x \in V \mid (s, \ell_1, x) \in E\}$ and $\{y \in V \mid (t, \overline{\ell}_1, y) \in E\}$ are non-empty, and whether there exists a path of even length from s to t in G. The first statements can be checked directly in FO, and the latter one can be checked in DynFO [13, 2]. This completes the proof of the first part of Theorem 1 in the case $n = 1$.

Remark 4. Note that this proof heavily relies on the property that the graph be undirected. In fact, the 1-letter Dyck reachability problem (over directed graphs) is bfo$^+$-reducible to the problem of computing distances in directed graphs, whose membership in DynFO or DynFO$^+$ is a long-standing open question [8, 2]. The reduction is as follows.

Given an unlabeled directed graph $G = (V, E)$, equip each edge with a label ℓ_1, and add self-loops (with the label ℓ_1) around each vertex in V. Then, for all vertices $v \in V$, add n vertices $(v, 1), \ldots, (v, n)$, where $n = |V|$, and add edges with the label $\overline{\ell}_1$ from v to $(v, 1)$ and from (v, i) to $(v, i + 1)$, for all i. It comes at once that the distance (in the original graph G) from a vertex s to a vertex t is k if and only if there exists a Dyck path (in the extended, labeled graph) from s to (t, k) but not to $(t, k - 1)$.

Furthermore, the proof of [8] showing that distances in graphs can be computed in DynTC^0 does not extend to the 1-letter Dyck reachability, whose precise dynamic complexity remains therefore unknown.

4 Two-letter Dyck reachability problem

We prove now that, for all integers $n \geq 2$, the n-letter Dyck reachability problem is P-complete for bfo$^+$ reductions.

We first introduce two auxiliary problems.

1. Let $G = (V, E)$ be an unlabeled directed graph, let $(V_\wedge, V_\vee)$ be a partition of V, and let s and t be two marked vertices of G. The *alternating reachability problem* asks whether s belongs to the smallest subset X of V such that all of $\{t\}$, $\{x \in V_\vee \mid \exists y \in X \text{ s.t. } (x, y) \in E\}$ and $\{x \in V_\wedge \mid \forall y \in V, (x, y) \in E \Rightarrow y \in X\}$ are subsets of X.

   Note that this problem could be alternatively and equivalently defined using the notion of winning state in a two-player turn-based zero-sum reachability game. However we choose the above definition using a fixed point to avoid defining the notion of winning strategies.

2. Let $G = (V, E, L)$ be a labeled directed graph with set of labels $L = V \cup \{\overline{v} \mid v \in V\} \cup \{\bullet\}$, where $\bullet$ is a fresh label symbol, and let s and t be two marked vertices of G. A *near-Dyck word* is an element of the set $\mathbf{D}'$ built over the grammar: $S \to \varepsilon \mid S \cdot \bullet \cdot S \mid v \cdot S \cdot \overline{v} \cdot S$ (for all $v \in V$). The *near-Dyck reachability problem* asks whether there exists a path π in G, with source s, sink t, and whose label belongs to $\mathbf{D}'$.

While it is well-known that the alternating reachability problem is P-hard for standard logarithmic-space reductions, it is also the case that it is P-hard for bfo$^+$ reductions [13]. Hence, we show in the two next subsections that there exists a bfo$^+$ reduction from the alternating reachability problem to the near-Dyck reachability problem, and that there exists a bfo$^+$ reduction from that latter problem to the 2-letter Dyck reachability problem. It will follow that the 2-letter (and, therefore, the n-letter) Dyck reachability problem is P-hard for bfo$^+$ reductions.

Remark 5. By carefully adaptating the reductions below, we might transform them into bfo reductions only. However, since the the alternating reachability problem is only known to be P-complete under bfo$^+$ reductions, such a transformation would not provide us with stronger P-completeness results for the n-letter Dyck reachability problem.

On the other hand, the n-letter Dyck reachability problem is clearly in P, as witnessed by the following algorithm, which will complete the proof of Theorem 1 in the case of the n-letter Dyck reachability problem.

Algorithm 1 Finding an n-letter Dyck path

$S \leftarrow \{(x, x) \mid x \in V\}$
while there exists pairs $(u, v) \notin S$ and $(u', v') \in S$ and edges (u, λ, u') and $(v', \overline{\lambda}, v)$ for some label $\lambda \in \{\ell_1, \ldots, \ell_n\}$ **do**
    $S \leftarrow S \cup \{(u, v)\}$
return $((s, t) \in S)$

4.1 From the near-Dyck reachability problem to the Dyck reachability problem

Let $G = (V, E, L)$ be a labeled directed graph with set of labels $L = V \cup \overline{V} \cup \{\bullet\}$ (where $\overline{V} = \{\overline{v} \mid v \in V\}$), and let s and t be two marked vertices of G.

We fix the new alphabet $\mathcal{L} = \{a, b, \overline{a}, \overline{b}\}$ and an arbitrary bijection $\theta : i \mapsto v_i$ between the set $\{1, \ldots, n\}$ and V (for $n = |V|$). We will expand the graph so that the letter $\bullet$ will be encoded using the word $a \cdot \overline{a}$, and each letter $v_i \in V$ (respectively, $\overline{v_i} \in \overline{V}$) will be represented by the word $a^i \cdot b \cdot a^{n+1-i}$ (respectively, $\overline{a}^{n+1-i} \cdot \overline{b} \cdot \overline{a}^i$).

Formally, let $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathcal{L})$ be the labeled directed graph defined by:

- $\mathcal{V} = V \cup \big(V \times \{\bullet\}\big) \cup \big(V \times (V \cup \overline{V}) \times \{0, 1, \ldots, n\}\big)$;
- $\mathcal{L} = \{a, b, \overline{a}, \overline{b}\}$;
- $\mathcal{E} = \mathcal{E}_1 \cup \mathcal{E}_2$, where

$$\begin{aligned}
\mathcal{E}_1 = {} & \{x \xrightarrow{a} (x, \bullet) \mid x \in V\} \cup \\
& \{x \xrightarrow{a} (x, v, 0) \mid x, v \in V\} \cup \\
& \{(x, v, i) \xrightarrow{a} (x, v, i+1) \mid x, v \in V, 0 \leq i \leq n-1, v \neq v_{i+1}\} \cup \\
& \{(x, v, i) \xrightarrow{b} (x, v, i+1) \mid x, v \in V, 0 \leq i \leq n-1, v = v_{i+1}\} \cup \\
& \{x \xrightarrow{\overline{a}} (x, \overline{v}, n) \mid x, v \in V\} \cup \\
& \{(x, \overline{v}, i+1) \xrightarrow{\overline{a}} (x, \overline{v}, i) \mid x, v \in V, 0 \leq i \leq n-1, v \neq v_{i+1}\} \cup \\
& \{(x, \overline{v}, i+1) \xrightarrow{\overline{b}} (x, \overline{v}, i) \mid x, v \in V, 0 \leq i \leq n-1, v = v_{i+1}\} \text{ and} \\
\mathcal{E}_2 = {} & \{(x, \bullet) \xrightarrow{\overline{a}} y \mid x \xrightarrow{\bullet} y \in E\} \cup \\
& \{(x, v, n) \xrightarrow{a} y \mid x \xrightarrow{v} y \in E\} \cup \\
& \{(x, \overline{v}, 0) \xrightarrow{\overline{a}} y \mid x \xrightarrow{\overline{v}} y \in E\},
\end{aligned}$$

and in which we mark the vertices s and t.

Each sequence of transitions $x \xrightarrow{a} (x, v_{i+1}, 0) \xrightarrow{a} \ldots \xrightarrow{a} (x, v_{i+1}, i) \xrightarrow{b} (x, v_{i+1}, i+1) \xrightarrow{a} \ldots \xrightarrow{a} \ldots (x, v_{i+1}, n)$ prepares the encoding of some edge leaving x with label v_{i+1}. If there is some edge $x \xrightarrow{v_{i+1}} y$ in the original graph, then only one edge $(x, v_{i+1}, n) \xrightarrow{a} y$ needs to be added: this is the role of the edges in $\mathcal{E}_2$. We use a similar encoding for edges labeled by $\overline{v_{i+1}}$, and an even simpler encoding for edges labeled by $\bullet$.

Proposition 6. *There exists a near-Dyck path from s to t in G if and only if there exists a Dyck path from s to t in $\mathcal{G}$.*

Proof. First, for every pair $(u, v) \in \mathcal{V}^2$, there exists at most one edge in $\mathcal{E}$ with source u and sink v. Henceforth, we omit representing labels of edges and of paths in $\mathcal{G}$.

We further define two mappings φ and ψ. The mapping φ identifies every label $\lambda \in L$ with a word $\varphi(\lambda) \in \mathcal{L}^*$, as follows:

$$\begin{aligned}
\varphi(\bullet) &= a \cdot \overline{a} \\
\varphi(v_i) &= a^i \cdot b \cdot a^{n+1-i} \text{ for all } i \in \{1, \ldots, n\} \\
\varphi(\overline{v_i}) &= \overline{a}^{n+1-i} \cdot \overline{b} \cdot \overline{a}^i \text{ for all } i \in \{1, \ldots, n\}
\end{aligned}$$

and extends immediately to a morphism from L^* to $\mathcal{L}^*$ that maps every near-Dyck word $w \in \mathbf{D}'$ to a Dyck word $\varphi(w) \in \mathbf{D}$. The mapping ψ identifies every edge $e \in E$ with a path $\psi(e)$ in $\mathcal{G}$, as follows:

$$\psi(x \xrightarrow{\bullet} y) = (x \to (x, \bullet) \to y)$$
$$\psi(x \xrightarrow{v} y) = (x \to (x, v, 0) \to \ldots \to (x, v, n) \to y) \text{ for all } v \in V$$
$$\psi(x \xrightarrow{\overline{v}} y) = (x \to (x, \overline{v}, n) \to \ldots \to (x, \overline{v}, 0) \to y) \text{ for all } \overline{v} \in \overline{V}$$

and extends immediately to a morphism that maps every path in G to a path in $\mathcal{G}$. The relation $\lambda(\psi(e)) = \varphi(\lambda(e))$ holds for all edges $e \in E$, and therefore extends to all paths π in G. Hence, a path π in G is near-Dyck if and only if the path $\psi(\pi)$ in $\mathcal{G}$ is Dyck.

In addition, let us call *nominal* paths in $\mathcal{G}$ the paths that belong to the set $\{\psi(e) \mid e \in E\}$, and *generic* paths in $\mathcal{G}$ the concatenations of nominal paths. Nominal paths are the minimal paths whose source and sink both belong to the subset V of $\mathcal{V}$. Hence, every path π from s to t in $\mathcal{G}$ is generic, thus π is the image by ψ of some path $\psi^{-1}(\pi)$ from s to t in G. □

The graph $\mathcal{G}$ is FO-definable as a function of G and of the bijection $\theta : \{1, \ldots, n\} \mapsto V$, and adding/deleting an edge in E amounts to adding/deleting exactly one edge in $\mathcal{E}_2$. Since θ can be precomputed in P, and due to Proposition 6, the near-Dyck reachability problem is therefore bfo^+-reducible to the Dyck reachability problem.

4.2 From the alternating reachability problem to the near-Dyck reachability problem

Let $G = (V, E)$ be an unlabeled directed graph, let $(V_\wedge, V_\vee)$ be a partition of V, and let s and t be two marked vertices of G.

We fix the new alphabet $\mathcal{L} = V \cup \{\overline{v} \mid v \in V\} \cup \{\bullet\}$, where $\bullet$ is a fresh label symbol, and an arbitrary bijection $\theta : i \mapsto v_i$ between the set $\{1, \ldots, n\}$ and V (for $n = |V|$). Recall the definition of the set X as a smallest fixed point page 5. We will expand the graph G so that the statement $s \in X$ be equivalent to the existence of some Dyck path from t to s (in the expanded graph). More precisely, for every vertex $x \in G$, we want the following statements to be equivalent to each other:

1. the vertex x belongs to X;
2. there exists a Dyck path from t to x (in the expanded graph);
3. there exists a path from t to itself whose label reduces to the letter x (in the expanded graph), where the reduction consists in deleting recursively the well-balanced subwords $\bullet$ and $v \cdot \overline{v}$ (for $v \in V$).

The equivalence $2 \Leftrightarrow 3$ is ensured by drawing an edge $x \xrightarrow{x} t$ for every vertex $x \in V$ and never using the label x for any other edge. The equivalence $1 \Leftrightarrow 2$ is obtained by induction, using the following construction (and the already proven equivalence $2 \Leftrightarrow 3$). For every vertex $x \in V_\vee$ and every edge $(x, y) \in E$, we draw an edge $y \xrightarrow{\bullet} x$; for every vertex $x \in V_\wedge$, we construct one unique path $\mathsf{in}(x)$ from t to x, and we build it so that its reduced label is $\prod_{y \in \{z \mid (x,z) \in E\}} \overline{v_y}$.

Formally, let $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathcal{L})$ be the labeled directed graph defined by:

- $\mathcal{V} = V \cup \big(V_\wedge \times \{0, 1, \ldots, n\}\big)$;
- $\mathcal{L} = V \cup \{\overline{v} \mid v \in V\} \cup \{\bullet\}$;
- $\mathcal{E} = \mathcal{E}_1 \cup \mathcal{E}_2$, where

$$\mathcal{E}_1 = \{x \xrightarrow{x} t \mid x \in V\} \cup \{t \xrightarrow{\bullet} (x, 0) \mid x \in V_\wedge\} \cup \{(x, n) \xrightarrow{\bullet} x \mid x \in V_\wedge\} \text{ and}$$
$$\mathcal{E}_2 = \{x \xrightarrow{\bullet} y \mid y \in V_\vee, (y, x) \in E\} \cup$$
$$\{(x, i) \xrightarrow{\overline{v_{i+1}}} (x, i+1) \mid x \in V_\wedge, 0 \leq i \leq n-1, (x, v_{i+1}) \in E\} \cup$$
$$\{(x, i) \xrightarrow{\bullet} (x, i+1) \mid x \in V_\wedge, 0 \leq i \leq n-1, (x, v_{i+1}) \notin E\},$$

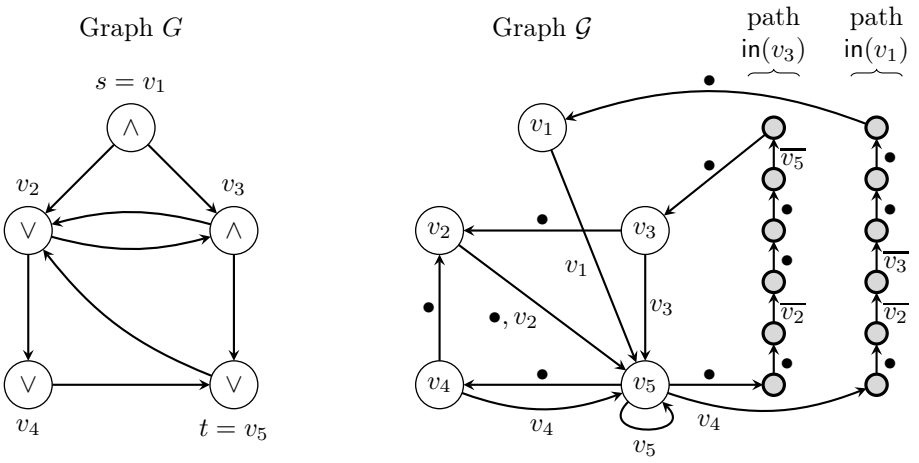

Figure 1: Graphs G and $\mathcal{G}$. A near-Dyck path witnessing that $s \in X$ is $v_5 \xrightarrow{v_5} v_5 \xrightarrow{\bullet} v_4 \xrightarrow{\bullet} v_2 \xrightarrow{v_2} v_5 \cdot \mathsf{in}(v_3) \cdot v_3 \xrightarrow{v_3} v_5 \xrightarrow{\bullet} v_4 \xrightarrow{\bullet} v_2 \xrightarrow{v_2} v_5 \cdot \mathsf{in}(v_1) \cdot v_1$

and in which we mark the vertices t and s.

The construction is illustrated in Fig. 1.

Proposition 7. *Let X be the smallest subset of V such that all of $\{t\}$, $\{x \in V_\vee \mid \exists y \in X \text{ s.t. } (x,y) \in E\}$ and $\{x \in V_\wedge \mid \forall y \in V, (x,y) \in E \Rightarrow y \in X\}$ are subsets of X. The vertex s belongs to X if and only if there exists a near-Dyck path from t to s in $\mathcal{G}$.*

Proof. Like in Section 4.1, observe that, for every pair $(u,v) \in \mathcal{V}^2$, there exists at most one edge in $\mathcal{E}$ with source u and sink v. Henceforth, we will sometimes omit representing labels of edges and of paths in $\mathcal{G}$. Conversely, for all $v \in V$, the edge $v \xrightarrow{v} t$ is the only edge in $\mathcal{E}$ with label v.

Furthermore, for all $x \in V_\wedge$, we denote by $\mathsf{in}(x)$ the path $t \to (x,0) \to (x,1) \to \ldots \to (x,n) \to x$. Every non-empty path with source in V and sink x must end with the sub-path $\mathsf{in}(x)$.

Observe that X is inductively defined as $X = \bigcup_{i \geq 0} X_i$, where $X_0 = \{t\}$ and

$$X_{i+1} = X_i \cup \{z \in V_\vee \mid \exists y \in X_i \text{ s.t. } (z,y) \in E\} \cup \{z \in V_\wedge \mid \forall y \in V, (z,y) \in E \Rightarrow y \in X_i\}.$$

Then, we say that a sequence $w_0, \ldots, w_k$ of vertices of G is *well-ordered* if $t = w_0$ and, for all $i \in \{1, \ldots, k\}$:

- if $w_i \in V_\vee$, then $\{z \in V \mid (w_i, z) \in E\} \cap \{w_0, \ldots, w_{i-1}\} \neq \emptyset$;
- if $w_i \in V_\wedge$, then $\{z \in V \mid (w_i, z) \in E\} \subseteq \{w_0, \ldots, w_{i-1}\}$.

For every $x \in X$, let $\iota(x)$ be the smallest i such that $x \in X_i$. Define any linear ordering $\prec$ over X such that $\iota(x) < \iota(y)$ implies $x \prec y$ (in particular the order between two vertices belonging to $X_i \setminus X_{i-1}$ is arbitrary). If $w_0 \prec w_1 \prec \ldots \prec w_k$ are the elements of X, with $k = |X| - 1$, then the sequence $w_0, \ldots, w_k$ is well-ordered. Consequently, every vertex $x \in X$ belongs to a well-ordered sequence, and we call *index* of x, which we denote by $\kappa(x)$, the smallest integer k such that x belongs to a well-ordered sequence $w_0, \ldots, w_k$. Conversely, if some vertex $x \in V$ belongs to a well-ordered sequence $w_0, \ldots, w_k$, then an immediate induction proves that $w_i \in X_i$ for all $i \leq k$, whence $x \in X$.

Consider now some vertex $x \in X$. We prove by induction on $\kappa(x)$ that there exits a near-Dyck path from t to x (in the graph $\mathcal{G}$). The result is immediate for $\kappa(x) = 0$, hence we assume that $\kappa(x) \geq 1$. Let $w_0, \ldots, w_k$ be a well-ordered sequence to which x belongs, with $k = \kappa(x)$. By minimality of $\kappa(x) = k$, we know that $x = w_k$. Whenever $0 \leq i \leq k-1$, the vertex w_i must belong to X, and its index is at most i, hence $\kappa(w_i) \leq i < \kappa(x)$. Hence, by induction hypothesis, there exists a near-Dyck path π^i from t to w_i in $\mathcal{G}$:

- If $x \in V_\vee$, there exists a vertex w_i (with $0 \leq i \leq k-1$) such that $(x, w_i) \in E$, hence the concatenation of π^i and of the one-edge path $w_i \xrightarrow{\bullet} x$ is a near-Dyck path from t to x.
- If $x \in V_\wedge$, let $w_1, \ldots, w_a$ be the elements of the set $\{z \in V \mid (x,z) \in E\}$ with $\theta^{-1}(w_1) < \ldots < \theta^{-1}(w_a)$. The concatenation of the paths $\pi^a, w_a \xrightarrow{w_a} t, \pi^{a-1}, \ldots, \pi^1, w_1 \xrightarrow{w_1} t$ and $\mathsf{in}(x)$ is a near-Dyck path from t to x (since $\mathsf{in}(x)$ is labeled by a word in $\bullet^+ \cdot \overline{w_1} \cdot \bullet^* \cdot \overline{w_2} \cdot \bullet^* \ldots \bullet^* \overline{w_a} \cdot \bullet^+$).

Conversely, let $x \in V$ be a vertex such that there exists a near-Dyck path from t to x in $\mathcal{G}$. Let π^x be a shortest such path, and let $l(x)$ be the length of π^x. We prove by induction on $l(x)$ that $x \in X$. The result is immediate for $l(x) = 0$, hence we assume that $l(x) \geq 1$:

- If $x \in V_\vee$, then π^x is the concatenation of some path ρ and of the edge $y \xrightarrow{\bullet} x$, for some $y \in V$. Hence ρ is a near-Dyck path of length $l(x)-1$, which proves that $l(y) < l(x)$, whence $y \in X$ (by induction hypothesis) and thus $x \in X$.
- If $x \in V_\wedge$, then π^x must end with $\mathsf{in}(x)$. The label of $\mathsf{in}(x)$ is a word belonging to $\bullet^+ \cdot \overline{w_1} \cdot \bullet^* \cdot \overline{w_2} \cdot \bullet^* \ldots \bullet^* \overline{w_a} \cdot \bullet^+$, where $\{w_1, w_2, \ldots, w_a\} = \{y \in V \mid (x,y) \in E\}$ and $\theta^{-1}(w_1) < \ldots < \theta^{-1}(w_a)$. Since $w_i \xrightarrow{w_i} t$ is the single edge labeled with w_i, we can write π^x as a concatenation of paths $\rho^{w_a}, w_a \xrightarrow{w_a} t, \rho^{w_{a-1}}, \ldots, \rho^{w_1}, w_1 \xrightarrow{w_1} t, \rho^t, \mathsf{in}(x)$, where every path ρ^z is a near-Dyck path from t to z (for $z \in V$). By induction (all such paths ρ^z have length smaller than $l(x)$), it follows immediately that all of $w_1, \ldots, w_k$ belong to X, whence $x \in X$.

□

The graph $\mathcal{G}$ is FO-definable as a function of G, of t and of the bijection $\theta : \{1, \ldots, n\} \mapsto V$, and adding/deleting an edge e in E amounts to adding/deleting exactly either one or two edges in $\mathcal{E}_2$. Since θ can be precomputed in P, and due to Proposition 7, the alternating reachability problem is therefore bfo$^+$-reducible to the near-Dyck reachability problem.

5 Two-letter undirected Dyck reachability problem

We proceed by proving that there exists a bfo$^+$ reduction from the 2-letter Dyck reachability problem (in directed graphs) to the 2-letter undirected reachability problem.

Let $G = (V, E, L)$ be a directed labeled graph, with $L = \{\ell_1, \ell_2, \overline{\ell}_1, \overline{\ell}_2\}$, and let s and t be two marked nodes of G. In addition, let $\mathcal{L} = \{0, 1, \overline{0}, \overline{1}\}$ be another set of labels.

The main difficulty, when working in an undirected graph, is that lots of cycles are created, generating lots of stuttering in the words labeling the path. It is therefore hard to really control where a path goes just by looking at its label. In particular, the recipe used in Section 4.1 to reduce the near-Dyck reachability problem to the 2-letter Dyck reachability problem cannot be used now, and it is not clear whether simple alternative reductions from the near-Dyck undirected reachability problem to the 2-letter undirected Dyck reachability problem exist. Hence, we prove directly the P-hardness of the 2-letter undirected Dyck reachability. In order to do so, we therefore rely on a rather intricate encoding, where each part plays an important role.

We denote by $\varphi : L^* \mapsto \mathcal{L}^*$ the homomorphism of monoids defined by:

$$\varphi(\ell_1) = 0 \cdot \overline{\mathbf{0}} \cdot \mathbf{1} \cdot 1 \cdot 0 \cdot 0 \cdot 1 \cdot 1 \cdot 1 \cdot 1 \cdot \overline{\mathbf{1}} \cdot \mathbf{0} \qquad \varphi(\overline{\ell}_1) = \overline{\mathbf{0}} \cdot \mathbf{1} \cdot \overline{1} \cdot \overline{1} \cdot \overline{1} \cdot \overline{1} \cdot \overline{0} \cdot \overline{0} \cdot \overline{1} \cdot \overline{\mathbf{1}} \cdot \mathbf{0} \cdot \overline{0}$$

$$\varphi(\ell_2) = 0 \cdot \overline{\mathbf{0}} \cdot \mathbf{1} \cdot 0 \cdot 0 \cdot 1 \cdot 1 \cdot 0 \cdot 0 \cdot 1 \cdot \overline{\mathbf{1}} \cdot \mathbf{0} \qquad \varphi(\overline{\ell}_2) = \overline{\mathbf{0}} \cdot \mathbf{1} \cdot \overline{1} \cdot \overline{0} \cdot \overline{0} \cdot \overline{1} \cdot \overline{1} \cdot \overline{0} \cdot \overline{0} \cdot \overline{\mathbf{1}} \cdot \mathbf{0} \cdot \overline{0}$$

Observe that the words $\varphi(\overline{\ell}_1)$ and $\varphi(\overline{\ell}_2)$ are formal inverses of the words $\varphi(\ell_1)$ and $\varphi(\ell_2)$.

In gray, we represent non-useful parts of the encoding: their only purpose is to provide equal-length encodings for ℓ_1, ℓ_2, $\overline{\ell}_1$ and $\overline{\ell}_2$. In red/bold are locks: *along a Dyck path*, once a lock has been traveled through, we cannot go back earlier in the encoding.

A last property of the encoding is the following. Every Dyck path traveling through the word $\varphi(\ell_1)$, may go back and forth arbitrarily, except at locks, which it may cross only once. Consider such a possible journey through the word $\varphi(\ell_1)$, and observe the word w obtained during that journey. This word is made of blocks that consist, alternatively, of letters 0 and $\overline{0}$, and of letters 1 and $\overline{1}$. Such a word, *even if we first reduce it* (by deleting recursively the words $0 \cdot \overline{0}$ and $1 \cdot \overline{1}$), will always satisfy the following properties:

1. there exists at least four such (non-empty) blocks;
2. the last block consists of letters 0 and $\overline{0}$;
3. the last two blocks are of odd length, and every other block is of even length.

To illustrate the above analysis, consider the direct journey through $\varphi(\ell_1)$, where we have identified the blocks: $(0\cdot\overline{0})\cdot(1\cdot 1)\cdot(0\cdot 0)\cdot(1\cdot 1\cdot 1\cdot 1\cdot\overline{1})\cdot(0)$. If that word is reduced, there remain four non-empty blocks: $(1\cdot 1)\cdot(0\cdot 0)\cdot(1\cdot 1\cdot 1)\cdot(0)$. One checks easily on this example that both words satisfy the properties 1–3.

Properties 1–2 hold for the word $\varphi(\ell_2)$, while property 3 should be replaced by:

3'. the first block of letters 1 and $\overline{1}$ and the last block of letters 0 and $\overline{0}$ are of odd length, and every other block is of even length.

This distinction between encodings of the two letters will allow identifying a path that encodes ℓ_1 or ℓ_2, even when there is backtracking between the two locks.

Finally, for every word $w \in \mathcal{L}^*$, we denote by w_i the i^{th} letter of w. Let $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathcal{L})$ be the undirected labeled graph defined by:

- $\mathcal{V} = V \cup V \times L \times V \times \{1, \dots, 11\}$;
- $\mathcal{E} = \mathcal{E}_{\mathsf{init}} \cup \mathcal{E}_{\mathsf{mid}} \cup \mathcal{E}_{\mathsf{end}}$, where

$$\mathcal{E}_{\mathsf{init}} = \{x \xleftrightarrow{\varphi(\lambda)_1} (x, \lambda, y, 1) \mid (x, \lambda, y) \in E\}$$
$$\mathcal{E}_{\mathsf{mid}} = \{(x, \lambda, y, i-1) \xleftrightarrow{\varphi(\lambda)_i} (x, \lambda, y, i) \mid (x, \lambda, y) \in E, 1 \leq i \leq 11\}$$
$$\mathcal{E}_{\mathsf{end}} = \{(x, \lambda, y, 11) \xleftrightarrow{\varphi(\lambda)_{12}} y \mid (x, \lambda, y) \in E\},$$

and in which we mark the vertices s and t.

Like in Sections 4.1 and 4.2, we observe an equivalence between the two kinds of Dyck reachability problems, as stated in the following result, whose proof is postponed to the Appendix.

Proposition 8. *There exists a Dyck path from s to t in G if and only if there exists a Dyck path from s to t in $\mathcal{G}$.*

The proof is more involved than the previous proofs, and will proceed as follows:

1. First, obviously, every Dyck path in G generates a Dyck path in $\mathcal{G}$ (via the morphism φ).
2. As already mentioned, *along a Dyck path* in $\mathcal{G}$, locks (in red/bold) will forbid long backtracking while reading encodings of letters; for instance, the first lock $\overline{\mathbf{0}}\cdot\mathbf{1}$ in $\varphi(\lambda)$ (with $\lambda \in \{\ell_1, \ell_2\}$) will forbid sub-paths of the form $(x, \lambda, y, 3) \xrightarrow{1} (x, \lambda, y, 2) \xrightarrow{\overline{0}} (x, \lambda, y, 1)$ since otherwise it would contradict the well-formedness of Dyck words.
3. We associate with every edge of G a class of paths in $\mathcal{G}$: an edge $e = (x, \lambda, y)$ is associated with the class $\mathcal{G}_e$ of those paths in $\mathcal{G}$ that consist in the concatenation of a (possibly empty) circuit around x with labels 0 and $\overline{0}$ only, then a path from x to y that uses all the vertices (x, λ, y, i) as internal vertices (and has no other internal vertex), and finally a circuit around y with labels 0 and $\overline{0}$ only. Then, we associate every path $\rho = (x_1, \lambda_1, y_1)\cdot\ldots\cdot(x_k, \lambda_k, y_k)$ in G with the class $\mathcal{G}_\rho = \mathcal{G}_{(x_1,\lambda_1,y_1)}\cdot\ldots\cdot\mathcal{G}_{(x_k,\lambda_k,y_k)}$ of paths in $\mathcal{G}$. It turns out that every Dyck path from s to t in $\mathcal{G}$ belongs to one unique class $\mathcal{G}_\rho$.
4. We focus on a notion of "approximate Dyck path": a path is approximately Dyck if its label is a sub-word of some Dyck word. Then, for every path ρ in G, we identify a sub-class of $\mathcal{G}$, which consists of those paths that belong to $\mathcal{G}_\rho$ and that are approximate Dyck paths. We show by induction that the latter sub-class is empty whenever ρ is not itself an approximate Dyck path, and that if $\mathcal{G}_\rho$ contains a Dyck path, then $\lambda(\rho)$ is a prefix of a Dyck word.
5. We also focus on an alternative notion of "well-balanced path" of a path in G: a path ρ with label $\lambda(\rho)$ is well-balanced if $\lambda(\rho)$ contains as many letters $\overline{\ell}_1$ or $\overline{\ell}_2$ as letters ℓ_1 or ℓ_2. We show that, if a class $\mathcal{G}_\rho$ contains a Dyck path, then ρ is well-balanced.

6. The conjunction of the steps 2–5 allows us to conclude that every Dyck path from s to t in $\mathcal{G}$ belongs to the class $\mathcal{G}_\rho$ of some Dyck path ρ from s to t in G.

The graph $\mathcal{G}$ is FO-definable as a function of G. Furthermore, adding/deleting an edge in E amounts to adding/deleting exactly twelve edges in $\mathcal{E}$. Due to Proposition 8, the 2-letter Dyck reachability problem is therefore bfo-reducible to the 2-letter undirected Dyck reachability problem.

On the other hand, as a restriction of the n-letter Dyck reachability problem, the n-letter undirected Dyck reachability problem is clearly in P, which completes the proof of Theorem 1.

References

[1] David A. Mix Barrington, Neil Immerman, and Howard Straubing. On uniformity within NC^1. *Journal of Computer and System Sciences*, 41(3):274–306, December 1990.

[2] Samir Datta, Raghav Kulkarni, Anish Mukherjee, Thomas Schwentick, and Thomas Zeume. Reachability is in DynFO. In *International Colloquium on Automata, Languages, and Programming*, pages 159–170. Springer, 2015.

[3] Stéphane Demri and Régis Gascon. The effects of bounding syntactic resources on Presburger LTL. *Journal of Logic and Computation*, 19(6):1541–1575, December 2009.

[4] Guozhu Dong and Jianwen Su. Incremental and decremental evaluation of transitive closure by first-order queries. *Information and Computation*, 120(1):101–106, 1995.

[5] Wouter Gelade, Marcel Marquardt, and Thomas Schwentick. The dynamic complexity of formal languages. *ACM Transactions on Computational Logic (TOCL)*, 13(3):19, 2012.

[6] Raymond Greenlaw, H. James Hoover, and Walter L. Ruzzo. *Limits to Parallel Computation: P-Completeness Theory*. Oxford University Press, 1995.

[7] Christoph Haase, Stephan Kreutzer, Joël Ouaknine, and James Worrell. Reachability in succinct and parametric one-counter automata. In *International Conference on Concurrency Theory*, pages 369–383. Springer, 2009.

[8] William Hesse. The dynamic complexity of transitive closure is in DynTC^0. *Theoretical Computer Science*, 296(3):473–485, 2003.

[9] William Hesse. *Dynamic Computational Complexity*. PhD thesis, Department of Computer Science, University of Massachusetts at Amherst, USA, September 2003.

[10] William Hesse and Neil Immerman. Complete problems for dynamic complexity classes. In *Proceedings of the 17th Annual Symposium on Logic in Computer Science (LICS'02)*, pages 313–322. IEEE Comp. Soc. Press, July 2002.

[11] Neil Immerman. *Descriptive complexity*. Graduate texts in computer science. Springer-Verlag, 1999.

[12] Peter Bro Miltersen, Sairam Subramanian, Jeffrey Scott Vitter, and Roberto Tamassia. Complexity models for incremental computation. *Theoretical Computer Science*, 130(1):203–236, 1994.

[13] Sushant Patnaik and Neil Immerman. Dyn-FO: A parallel, dynamic complexity class. *Journal of Computer and System Sciences*, 55(2):199–209, October 1997.

[14] Volker Weber and Thomas Schwentick. Dynamic complexity theory revisited. *Theory of Computing Systems*, 40(4):355–377, June 2007.

— Appendix: Proving Proposition 8 —

For every pair $(u, v) \in \mathcal{V}^2$ there exists at most one undirected edge between u and v in $\mathcal{E}$. Henceforth, we may omit representing labels of edges and of paths in $\mathcal{G}$.

A first result comes quickly.

Proposition 9. *Let ρ be a Dyck path from s to t in G. There exists a Dyck path from s to t in $\mathcal{G}$.*

Proof. Let us denote by ψ the mapping that identifies every edge $e \in E$ with the path $\psi(e) = (x \to (x, \theta, y, 1) \to \ldots \to (x, \theta, y, 11) \to y)$ in $\mathcal{G}$, where $e = (x, \theta, y)$. Observe that ψ extends immediately to a morphism that maps every path in G to a path in $\mathcal{G}$. The relation $\lambda(\psi(e)) = \varphi(\lambda(e))$ holds for all edges $e \in E$, and therefore extends to all paths in G. Hence, a path ρ in G is Dyck if and only if the path $\psi(\rho)$ in $\mathcal{G}$ is Dyck. □

However, unlike in Section 4.1, there may exist Dyck paths in $\mathcal{G}$ that are not of the form $\psi(\rho)$, as shown by the example of the Dyck cycle $s_1 \xrightarrow{0} (s_1, \ell_1, s_2, 1) \xrightarrow{0} s_1 \xrightarrow{\overline{0}} (s_2, \overline{\ell}_2, s_1, 11) \xrightarrow{\overline{0}} s_1$ in the graph displayed in Fig. 2.

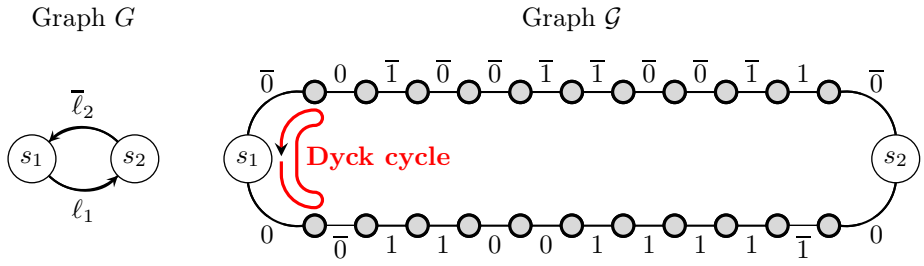

Figure 2: Graphs G and $\mathcal{G}$, and Dyck cycle in $\mathcal{G}$

Consequently, we cannot use directly the morphism ψ to associate every Dyck path in $\mathcal{G}$ with a Dyck path in G.

We overcome this problem as follows. Let $\mathcal{Q}$ be the set of all sub-words of all Dyck words with letters in $\mathcal{L}$ (called *approximate Dyck words*), and let $\mathcal{P}$ the set of all paths π in $\mathcal{G}$ such that $\lambda(\pi) \in \mathcal{Q}$ (called *approximate Dyck paths*). Moreover, for every set S of paths, we denote by $\lambda(S)$ the set of labels of paths in S, i.e. $\lambda(S) = \{\lambda(\pi) \mid \pi \in S\}$. It comes at once that $\lambda(\mathcal{P}) \subseteq \mathcal{Q}$, that $\mathcal{Q}$ is sub-word closed, and that none of the words $1 \cdot \overline{0}$ nor $0 \cdot \overline{1}$ belongs to $\mathcal{Q}$. It follows that no *doomed* path belongs to $\mathcal{P}$, where the notion of doomed path is defined below. This was the role of the locks mentioned in the core of the paper.

Definition 10. A 2-letter path π in $\mathcal{G}$ is said to be *doomed* if there exists an edge $(x, \lambda, y) \in E$ such that either

- $\lambda \in \{\ell_1, \ell_2\}$ and $\pi = (x, \lambda, y, 3) \xrightarrow{1} (x, \lambda, y, 2) \xrightarrow{\overline{0}} (x, \lambda, y, 1)$ or $\pi = y \xrightarrow{0} (x, \lambda, y, 12) \xrightarrow{\overline{1}} (x, \lambda, y, 11)$;
- $\lambda \in \{\overline{\ell}_1, \overline{\ell}_2\}$ and $\pi = (x, \lambda, y, 2) \xrightarrow{1} (x, \lambda, y, 1) \xrightarrow{\overline{0}} x$ or $\pi = (x, \lambda, y, 12) \xrightarrow{0} (x, \lambda, y, 11) \xrightarrow{\overline{1}} (x, \lambda, y, 10)$.

We further say that a path in $\mathcal{G}$ is a *nominal* path if it is a minimal path whose source and sink belong to V. For all edges $(x, \lambda, y) \in E$, we denote by $\mathcal{P}_{x,\lambda,y}$ the set of nominal paths $\pi \in \mathcal{P}$ such that π has source x, sink y, and such that its internal vertices are exactly the elements of the set $\{(x, \lambda, y, i) \mid 1 \leq i \leq 11\}$. For all vertices $x \in V$, we also denote by $\mathcal{P}_x$ the set of nominal paths $\pi \in \mathcal{P}$ such that π has source and sink x, and whose edges are all labeled with 0 or $\overline{0}$. These two classes of paths capture the entire family of nominal paths that belong to $\mathcal{P}$, as shown by the following result.

Lemma 11. *Let $\pi \in \mathcal{P}$ be a nominal path in $\mathcal{G}$. Either there exists an edge $(x, \lambda, y) \in E$ such that $\pi \in \mathcal{P}_{x,\lambda,y}$ or there exists a vertex $x \in V$ such that $\pi \in \mathcal{P}_x$. Moreover, the sets $\mathcal{P}_{x,\lambda,y}$ and $\mathcal{P}_x$ are pairwise disjoint.*

Proof. We first assume that some edge e in π is labeled by 1 or $\overline{1}$. By construction, there exists a unique edge $(x, \lambda, y) \in E$ and a unique pair of integers $i, j \in \{1, \dots, 11\}$ such that $e = (x, \lambda, y, i) \to (x, \lambda, y, j)$, with $j = i \pm 1$. Since π is nominal, its internal vertices belong to the set $\{(x, \lambda, y, i) \mid 1 \leq i \leq 11\}$, and its source and sink belong to $\{x, y\}$. Then, since π belongs to $\mathcal{P}$, it does not contain doomed paths, hence its source must be x and its sink must be y.

Then, assume that no edge in π is labeled by 1 or $\overline{1}$. Since π is nominal, there exists a unique edge $(x, \lambda, y) \in E$ such that the internal vertices of π belong to the set $\{(x, \lambda, y, i) \mid 1 \leq i \leq 11\}$, and its source and sink belong to $\{x, y\}$. If x is the source of π, then π can never reach the vertex $(x, \lambda, y, 3)$, hence x is the sink of π; if y is the source of π, then π can never reach the vertex $(x, \lambda, y, 9)$, hence y is the sink of π.

Observing that every path in every set $\mathcal{P}_{x,\lambda,y}$ contains an edge labeled by 1 or $\overline{1}$ completes the proof. □

Going further, we associate with every path $\rho = (v_1, \lambda_1, w_1) \cdot \ldots \cdot (v_k, \lambda_k, w_k)$ in G the set $\overline{\mathcal{P}}_\rho$ of paths in $\mathcal{G}$ defined by:

$$\overline{\mathcal{P}}_\rho = \mathcal{P}^*_{v_1} \cdot \mathcal{P}_{v_1,\lambda_1,w_1} \cdot \mathcal{P}^*_{v_2} \cdot \mathcal{P}_{v_2,\lambda_2,w_2} \cdot \ldots \cdot \mathcal{P}^*_{v_k} \cdot \mathcal{P}_{v_k,\lambda_k,w_k} \cdot \mathcal{P}^*_{w_k}.$$

Observe that, unlike the sets $\mathcal{P}_x$ and $\mathcal{P}_{x,\lambda,y}$, the sets $\overline{\mathcal{P}}_\rho$ may contain paths that are not nominal and/or not approximate Dyck paths.

Definition 12. Let π be a Dyck path in $\mathcal{G}$ from s to t. There exists a unique sequence of vertices $v_0, \dots, v_k$, a unique partial function $f_\pi : \{1, \dots, k\} \mapsto L$ and a unique sequence of nominal paths $\pi_1, \dots, \pi_k$ such that:

- $v_0 = s$ and $v_k = t$;
- for all $i \in \mathsf{dom}(f_\pi)$, the edge $(v_{i-1}, f_\pi(i), v_i)$ belongs to E, and $\pi_i \in \mathcal{P}_{v_{i-1},f_\pi(i),v_i}$;
- for all $i \in \{1, \dots, k\} \setminus \mathsf{dom}(f_\pi)$, we have $v_{i-1} = v_i$, and $\pi_i \in \mathcal{P}_{v_i}$;
- $\pi = \pi_1 \cdot \ldots \cdot \pi_k$.

We call *nominal vertex sequence* of π sequence $v_0, \dots, v_k$, *nominal label mapping* of π the mapping f_π, *nominal decomposition* of π the sequence $\pi_1, \dots, \pi_k$, and *nominal ancestor* of π the path $(v_{i-1}, f_\pi(i), v_i)_{i \in \mathsf{dom}(f_\pi)}$.

It comes immediately that, for every Dyck path π in $\mathcal{G}$, the nominal ancestor of π is the unique path ρ in G such that $\pi \in \overline{\mathcal{P}}_\rho$.

Now, we denote by $\mathcal{Q}_{\mathsf{init}}$ the set of all prefixes of all Dyck words with letters in $\mathcal{L}$. Observe that, for all words $\rho_1, \rho_2 \in \mathcal{L}^*$, the three words $\rho_1 \cdot \rho_2$, $\rho_1 \cdot 0 \cdot \overline{0} \cdot \rho_2$ and $\rho_1 \cdot 1 \cdot \overline{1} \cdot \rho_2$ are either all Dyck or all non-Dyck.

Then, for every word $w \in \mathcal{L}^*$, we call *reduced word* of w the word $\mathsf{red}(w)$ obtained from w by deleting recursively the 2-letter words $0 \cdot \overline{0}$ or $1 \cdot \overline{1}$. Alternatively, if we consider w as an element of the free group generated by ℓ_1 and ℓ_2 (with inverses $\overline{\ell}_1$ and $\overline{\ell}_2$), then $\mathsf{red}(w)$ is the reduced word representing w.

We just proved that w is Dyck if and only if $\mathsf{red}(w)$ is Dyck. Henceforth, we identify regular expressions with the language they represent and, for all sets of words A and B, we write $A \overset{\mathsf{red}}{\subseteq} B$ as a placeholder for $\{\mathsf{red}(w) \mid w \in A\} \subseteq B$.

Recall that no word in $\mathcal{Q}$ contains the sub-words $1 \cdot \overline{0}$ nor $0 \cdot \overline{1}$, and that no word whose first letter is $\overline{0}$ or $\overline{1}$ can belong to $\mathcal{Q}_{\mathsf{init}}$. Hence, in a reduced word $w \in \mathcal{Q}$, no letter 0 or 1 can be followed by a letter $\overline{0}$ or $\overline{1}$ and, if $w \in \mathcal{Q}_{\mathsf{init}}$, then w contains no letter $\overline{0}$ or $\overline{1}$.

In addition, consider the following regular expressions:

$$\omega_+ = (0 \cdot 0 + 1 \cdot 1)^* \qquad \varpi_+ = (\omega_+ \cdot 1 \cdot \omega_+ \cdot 1)^* \cdot \omega_+$$
$$\omega_- = (\overline{0} \cdot \overline{0} + \overline{1} \cdot \overline{1})^* \qquad \varpi_- = (\omega_- \cdot \overline{1} \cdot \omega_- \cdot \overline{1})^* \cdot \omega_-$$
$$\omega = (\omega_+ + \omega_- + \overline{0} \cdot 0)^* \qquad \varpi = (\omega \cdot 1 \cdot \omega \cdot \overline{1})^* \cdot \omega$$

The intuition behind these expressions follows the sketch of proof of Proposition 8, in particular the notion of approximate Dyck path. Roughly speaking, and anticipating Lemma 15, we prove below that:

- every circuit in $\mathcal{G}$ whose edges are all labeled by 0 or $\overline{0}$ has its label in ω, and therefore in ϖ;
- if ρ is an approximate Dyck path in G, every approximate Dyck path in $\mathcal{P}_\rho$ has its label in ϖ.

These statements are proved by induction, hence ϖ was found as a least-fixed point for some closure properties similar to those of Dyck paths. Then, ω_+, ω_-, ϖ_+ and ϖ_- are just specializations of ω and ϖ where only labels 0 or 1 (respectively, $\overline{0}$ or $\overline{1}$) are allowed.

These latter regular expressions will allow us to express the following auxiliary results, which will lead us to Lemma 15.

Lemma 13. *The following relations hold:*

$$(1 \cdot 0 \cdot \varpi) \cap \mathcal{Q} \overset{\mathsf{red}}{\subseteq} 1 \cdot 0 \cdot \varpi_+ \qquad\qquad (\varpi \cdot \overline{0} \cdot \overline{1}) \cap \mathcal{Q} \overset{\mathsf{red}}{\subseteq} \varpi_- \cdot \overline{0} \cdot \overline{1}$$

Proof. Consider the regular expression $\alpha = 1 \cdot (0+\overline{0}) \cdot (0 \cdot 0 + \overline{0} \cdot 0 + 0 \cdot \overline{0} + \overline{0} \cdot \overline{0} + 1 + \overline{1})^*$. It comes at once that $1 \cdot 0 \cdot \varpi \subseteq \alpha$. Furthermore, for all words ρ_1, ρ_2 such that either $\rho_1 \cdot 0 \cdot \overline{0} \cdot \rho_2$ or $\rho_1 \cdot 1 \cdot \overline{1} \cdot \rho_2$ belongs to α, the word $\rho_1 \cdot \rho_2$ belongs to α too, which means that α is closed under reduction.

Let ρ be a word in ϖ such that $1 \cdot 0 \cdot \rho \in \mathcal{Q}$, and let $\rho' = \mathsf{red}(1 \cdot 0 \cdot \rho)$. Since ρ belongs to α, so does ρ', whose leftmost letter must then be 1. We remarked above that, since $\rho' \in \mathcal{Q}$, no letter 0 or 1 in ρ' can be followed by a letter $\overline{0}$ or $\overline{1}$. It follows that ρ' does not contain any letter $\overline{0}$ or $\overline{1}$, i.e. that $\rho' = 1 \cdot 0 \cdot \rho''$ for some reduced word ρ''. Then, ρ'' must be the reduction of ρ, and cannot contain letters $\overline{0}$ or $\overline{1}$, hence $\rho'' \in \varpi_+$, i.e. $\rho' \in 1 \cdot 0 \cdot \varpi_+$. We prove similarly the relation $(\varpi \cdot \overline{0} \cdot \overline{1}) \cap \mathcal{Q} \overset{\mathsf{red}}{\subseteq} \varpi_- \cdot \overline{0} \cdot \overline{1}$. □

Lemma 14. *For all vertices $x, y \in V$, we have $\lambda(\mathcal{P}_x) \overset{\mathsf{red}}{\subseteq} \varpi$ and*

$$\lambda(\mathcal{P}_{x,\ell_1,y}) \overset{\mathsf{red}}{\subseteq} \varpi \cdot 1 \cdot 1 \cdot 0 \cdot 0 \cdot \omega_+ \cdot 1 \cdot 0 \qquad \lambda(\mathcal{P}_{x,\ell_2,y}) \overset{\mathsf{red}}{\subseteq} \varpi \cdot 1 \cdot \omega_+ \cdot 0 \cdot 0 \cdot 1 \cdot 1 \cdot \omega_+ \cdot 0$$
$$\lambda(\mathcal{P}_{x,\overline{\ell}_1,y}) \overset{\mathsf{red}}{\subseteq} \overline{0} \cdot \overline{1} \cdot \omega_- \cdot \overline{0} \cdot \overline{0} \cdot \overline{1} \cdot \overline{1} \cdot \varpi \qquad \lambda(\mathcal{P}_{x,\overline{\ell}_2,y}) \overset{\mathsf{red}}{\subseteq} \overline{0} \cdot \omega_- \cdot \overline{1} \cdot \overline{1} \cdot \overline{0} \cdot \overline{0} \cdot \omega_- \cdot \overline{1} \cdot \varpi$$

Proof. First, let π be a path in $\mathcal{P}_x$. By construction, the graph $\mathcal{G}$ is bipartite, hence π is of even length. Its edges are labeled by 0 and $\overline{0}$ only, hence the word $\lambda(\pi)$, once reduced, belongs to ω, and therefore to ϖ.

Then, let π be a path in $\mathcal{P}_{x,\ell_1,y}$. By construction, π belongs to $\mathcal{P}$, hence its label $\lambda(\pi)$ belongs to $\mathcal{Q}$, and π cannot contain any doomed path. Consequently, we can decompose π as a concatenation of paths $\pi_1 \cdot \pi_2 \cdot \pi_3 \cdot \pi_4$, where:

- π_1 goes from x to $(x, \ell_1, y, 2)$, with internal vertices in $\{(x, \ell_1, y, i) \mid 1 \leq i \leq 2\}$;
- π_2 goes from $(x, \ell_1, y, 2)$ to $(x, \ell_1, y, 6)$, with internal vertices in $\{(x, \ell_1, y, i) \mid 2 \leq i \leq 6\}$;
- π_3 goes from $(x, \ell_1, y, 6)$ to $(x, \ell_1, y, 11)$, with internal vertices in $\{(x, \ell_1, y, i) \mid 2 \leq i \leq 11\}$;
- π_4 is the one-letter path $(x, \ell_1, y, 11) \to y$ (by definition of a nominal path, we stop at the first occurrence of y).

It comes at once that $\lambda(\pi_1) \in \varpi$, that $\lambda(\pi_2) \in 1 \cdot 1 \cdot \omega_+ \cdot 0 \cdot 0$ and that $\lambda(\pi_4) = 0$. The word $\lambda(\pi_3)$ belongs a priori to $\omega \cdot \overline{1}$, but we show that when reduced, it actually contains no occurrence of letters $\overline{0}$ and $\overline{1}$. The case of $\overline{0}$ is obvious.

Assume now that there is some $\overline{1}$ in the reduced version of $\lambda(\pi_3)$. The leftmost letter $\overline{1}$ cannot be preceded by a 1 (since $\lambda(\pi_3)$ is reduced) nor by a 0 (since $0 \cdot \overline{1} \notin \mathcal{Q}$). Hence, it must be the first letter of $\lambda(\pi_3)$. However, since 0 is the last letter of $\lambda(\pi_2)$, the concatenated word

$\lambda(\pi) = \lambda(\pi_1)\cdot\lambda(\pi_2)\cdot\lambda(\pi_3)\cdot\lambda(\pi_4)$ cannot belong to $\mathcal{Q}$ itself, which is a contradiction. We conclude that the reduced version of $\lambda(\pi_3)$ belongs to $\omega_+ \cdot 1$. Hence, we have proved that

$$\lambda(\mathcal{P}_{x,\ell_1,y}) \overset{\mathsf{red}}{\subseteq} \varpi \cdot 1 \cdot 1 \cdot \omega_+ \cdot 0 \cdot 0 \cdot \omega_+ \cdot 1 \cdot 0 = \varpi \cdot 1 \cdot 1 \cdot 0 \cdot 0 \cdot \omega_+ \cdot 1 \cdot 0.$$

The case of the language $\lambda(\mathcal{P}_{x,\overline{\ell}_1,y}) \cap \mathcal{Q}$ is treated in the exact same manner, and the two remaining cases are very similar as well. □

Using the two previous technical lemmas, we are able to show that the set $\mathcal{Q}$ of *approximate Dyck words* allows to discriminate between legal (for Dyck) and illegal consecutions of nominal paths. Similarly, prefixes $\mathcal{Q}_{\mathsf{init}}$ eliminate illegal beginnings of Dyck paths. This is stated as follows.

Lemma 15. *For all vertices $x, x', y, y' \in V$, we have*

$$(\lambda(\mathcal{P}_{x,\ell_1,y}) \cdot \varpi \cdot \lambda(\mathcal{P}_{x',\overline{\ell}_1,y'})) \cap \mathcal{Q} \overset{\mathsf{red}}{\subseteq} \varpi \qquad (\lambda(\mathcal{P}_{x,\ell_2,y}) \cdot \varpi \cdot \lambda(\mathcal{P}_{x',\overline{\ell}_2,y'})) \cap \mathcal{Q} \overset{\mathsf{red}}{\subseteq} \varpi$$

$$(\lambda(\mathcal{P}_{x,\ell_1,y}) \cdot \varpi \cdot \lambda(\mathcal{P}_{x',\overline{\ell}_2,y'})) \cap \mathcal{Q} = \emptyset \qquad (\lambda(\mathcal{P}_{x,\ell_2,y}) \cdot \varpi \cdot \lambda(\mathcal{P}_{x',\overline{\ell}_1,y'})) \cap \mathcal{Q} = \emptyset$$

$$(\varpi \cdot \lambda(\mathcal{P}_{x,\overline{\ell}_1,y})) \cap \mathcal{Q}_{\mathsf{init}} = \emptyset \qquad (\varpi \cdot \lambda(\mathcal{P}_{x,\overline{\ell}_2,y})) \cap \mathcal{Q}_{\mathsf{init}} = \emptyset$$

Proof. First, let ρ_1, ρ_2 and ρ_3 be reduced words in $\lambda(\mathcal{P}_{x,\ell_1,y})$, ϖ and $\lambda(\mathcal{P}_{x',\overline{\ell}_1,y'})$ such that $\rho_1\cdot\rho_2\cdot\rho_3$ belongs to $\mathcal{Q}$. Lemma 14 proves that ρ_1 ends with the suffix $1\cdot 0$, and that ρ_3 begins with the prefix $\overline{0}\cdot\overline{1}$.

Hence, the word $\rho_2' = 1\cdot 0\cdot\rho_2\cdot\overline{0}\cdot\overline{1}$ belongs to $\mathcal{Q}$, and therefore Lemma 13 proves that ρ_2' belongs to both $1\cdot 0\cdot\varpi_+\cdot\overline{0}\cdot\overline{1}$ and $1\cdot 0\cdot\varpi_-\cdot\overline{0}\cdot\overline{1}$. Since $(1\cdot 0\cdot\varpi_+\cdot\overline{0}\cdot\overline{1})\cap(1\cdot 0\cdot\varpi_-\cdot\overline{0}\cdot\overline{1}) \overset{\mathsf{red}}{\subseteq} 1\cdot 0\cdot\overline{0}\cdot\overline{1} \overset{\mathsf{red}}{\subseteq} \varepsilon$, it follows that $\rho_1\cdot\rho_2\cdot\rho_3$, once reduced, belongs to the set $\varpi\cdot 1\cdot 1\cdot 0\cdot 0\cdot\omega_+\cdot\varepsilon\cdot\omega_-\cdot\overline{0}\cdot\overline{0}\cdot\overline{1}\cdot\overline{1}\cdot\varpi$, i.e. to ϖ itself.

We prove the relation $(\lambda(\mathcal{P}_{x,\ell_2,y})\cdot\varpi\cdot\lambda(\mathcal{P}_{x',\overline{\ell}_2,y'}))\cap\mathcal{Q} \overset{\mathsf{red}}{\subseteq} \varpi$ in a similar way. Let ρ_1, ρ_2 and ρ_3 be reduced words in $\lambda(\mathcal{P}_{x,\ell_2,y})$, ϖ and $\lambda(\mathcal{P}_{x',\overline{\ell}_2,y'})$ such that $\rho_1\cdot\rho_2\cdot\rho_3$ belongs to $\mathcal{Q}$. Lemma 14 proves that ρ_1 belongs to $\varpi\cdot 1\cdot\omega_+\cdot 0\cdot 0\cdot\omega_+\cdot 1\cdot 1\cdot 0^{2a+1}$ and that ρ_3 belongs to $\overline{0}^{2b+1}\cdot\overline{1}\cdot\overline{1}\cdot\omega_-\cdot\overline{0}\cdot\overline{0}\cdot\omega_-\cdot\overline{1}\cdot\varpi$ for some integers $a, b \geq 0$.

Once again, Lemma 13 proves that the word $\rho_2' = 1\cdot 0^{2a+1}\cdot\rho_2\cdot\overline{0}^{2b+1}\cdot\overline{1}$ is reducible to the empty word, which shows that $\rho_1\cdot\rho_2\cdot\rho_3$, once reduced, belongs to $\varpi\cdot 1\cdot\omega_+\cdot 0\cdot 0\cdot\omega_+\cdot 1\cdot\varepsilon\cdot\overline{1}\cdot\omega_-\cdot\overline{0}\cdot\overline{0}\cdot\omega_-\cdot\overline{1}\cdot\varpi$, i.e. to ϖ itself.

Second, assume that there exists reduced words ρ_1, ρ_2 and ρ_3 in $\lambda(\mathcal{P}_{x,\ell_1,y})$, ϖ and $\lambda(\mathcal{P}_{x',\overline{\ell}_2,y'})$ such that $\rho_1\cdot\rho_2\cdot\rho_3$ belongs to $\mathcal{Q}$. The same arguments as above show that $\rho_1\cdot\rho_2\cdot\rho_3$, once reduced, belongs to $\varpi\cdot 1\cdot 1\cdot 0\cdot 0\cdot\omega_+\cdot\varepsilon\cdot\overline{1}\cdot\omega_-\cdot\overline{0}\cdot\overline{0}\cdot\omega_-\cdot\overline{1}\cdot\varpi$, i.e. that there exists integers $a, b \geq 0$ such that $\mathsf{red}(\rho_1\cdot\rho_2\cdot\rho_3)$ belongs to $\varpi\cdot 1\cdot 1\cdot\omega_+\cdot 0\cdot 0\cdot 1^{2a}\cdot\varepsilon\cdot\overline{1}^{2b+1}\cdot\overline{0}\cdot\overline{0}\cdot\omega_-\cdot\omega_-\cdot\overline{1}\cdot\varpi$.

Every word of the form $0\cdot 1^{2a}\cdot\overline{1}^{2b+1}\cdot\overline{0}$ is reducible to a word $0\cdot 1^{2x+1}\cdot\overline{0}$ or $0\cdot\overline{1}^{2x+1}\cdot\overline{0}$, hence it cannot belong to $\mathcal{Q}$. This proves that $\rho_1\cdot\rho_2\cdot\rho_3$ cannot belong to $\mathcal{Q}$ as well.

We prove in the exact same manner the equality $(\lambda(\mathcal{P}_{x,\ell_2,y})\cdot\varpi\cdot\lambda(\mathcal{P}_{x',\overline{\ell}_1,y'}))\cap\mathcal{Q} = \emptyset$.

Finally, the same arguments also show that $(\varpi\cdot\lambda(\mathcal{P}_{x,\overline{\ell}_1,y}))\cap Q \overset{\mathsf{red}}{\subseteq} \varpi_-\cdot\lambda(\mathcal{P}_{x,\overline{\ell}_1,y})$ and that $(\varpi\cdot\lambda(\mathcal{P}_{x,\overline{\ell}_2,y}))\cap Q \overset{\mathsf{red}}{\subseteq} \varpi_-\cdot\lambda(\mathcal{P}_{x,\overline{\ell}_2,y})$. However, no word in $\mathcal{Q}_{\mathsf{init}}$ begins with a letter $\overline{0}$ or $\overline{1}$, while every word in ϖ_- (if that word is non-empty), in $\lambda(\mathcal{P}_{x,\overline{\ell}_1,y})$ or in $\lambda(\mathcal{P}_{x,\overline{\ell}_2,y})$ begins with such a letter. Hence, both$(\varpi\cdot\lambda(\mathcal{P}_{x,\overline{\ell}_1,y}))\cap Q$ and $(\varpi\cdot\lambda(\mathcal{P}_{x,\overline{\ell}_2,y}))\cap Q$ are empty. □

This allows to state the partial correctness of the construction.

Proposition 16. *Let ρ be a path in G. If ρ is a Dyck path, then $\lambda(\overline{\mathcal{P}}_\rho)\cap\mathcal{Q} \overset{\mathsf{red}}{\subseteq} \varpi$. If ρ is not an approximate Dyck path, then $\lambda(\overline{\mathcal{P}}_\rho)\cap\mathcal{Q} = \emptyset$. Finally, if $\lambda(\rho)$ is not a prefix of a Dyck word, then $\lambda(\overline{\mathcal{P}}_\rho)\cap\mathcal{Q}_{\mathsf{init}} = \emptyset$.*

Proof. We first prove that $\lambda(\overline{\mathcal{P}}_\rho)\cap\mathcal{Q} \overset{\mathsf{red}}{\subseteq} \varpi$ if ρ is a Dyck path, and proceed by induction on the length $|\rho|$ of the path ρ. If $|\rho| = 0$, then $\lambda(\overline{\mathcal{P}}_\rho) \subseteq \lambda(\mathcal{P}_x)^* \subseteq \varpi^* = \varpi$, where x is the source and sink of ρ. Then, if $|\rho| \geq 1$ and if ρ is a concatenation of two non-empty Dyck paths ρ_1 and ρ_2, we have

$$\lambda(\overline{\mathcal{P}}_\rho)\cap\mathcal{Q} \overset{\mathsf{red}}{\subseteq} (\lambda(\overline{\mathcal{P}}_{\rho_1})\cap\mathcal{Q})\cdot(\lambda(\overline{\mathcal{P}}_{\rho_2})\cap\mathcal{Q}) \overset{\mathsf{red}}{\subseteq} \varpi\cdot\varpi = \varpi.$$

Finally, if $|\rho| \geq 1$ and if ρ is a concatenation of an edge (x_1, θ, x_2), a Dyck path ρ' with source x_2 and sink y_1, and an edge $(y_1, \overline{\theta}, y_2)$ (with $\theta \in \{\ell_1, \ell_2\}$), then

$$\begin{aligned}\lambda(\overline{\mathcal{P}}_\rho) \cap \mathcal{Q} &\overset{\mathsf{red}}{\subseteq} \lambda(\mathcal{P}_{x_1})^* \cdot ((\lambda(\mathcal{P}_{x_1,\theta,x_2}) \cdot (\lambda(\overline{\mathcal{P}}_{\rho'}) \cap \mathcal{Q}) \cdot \lambda(\mathcal{P}_{y_1,\overline{\theta},y_2})) \cap \mathcal{Q}) \cdot \lambda(\mathcal{P}_{y_2})^* \\ &\overset{\mathsf{red}}{\subseteq} \varpi^* \cdot ((\lambda(\mathcal{P}_{x_1,\theta,x_2}) \cdot \varpi \cdot \lambda(\mathcal{P}_{y_1,\overline{\theta},y_2})) \cap \mathcal{Q}) \cdot \varpi^* \\ &\overset{\mathsf{red}}{\subseteq} \varpi^* \cdot \varpi \cdot \varpi^* = \varpi.\end{aligned}$$

Then, we prove that $\lambda(\overline{\mathcal{P}}_\rho) \cap \mathcal{Q} = \emptyset$ if ρ is not an approximate Dyck path. Indeed, let ρ' be a minimal sub-path of ρ' that is not approximate Dyck. The word ρ' cannot be empty, hence it must be of the form $\rho' = (x_1, \theta, x_2) \cdot \rho'' \cdot (y_1, \chi, y_2)$, where ρ'' is a Dyck path and either $\theta = \ell_1$ and $\chi = \overline{\ell}_2$ or $\theta = \ell_2$ and $\overline{\chi} = \overline{\ell}_1$. It follows that

$$\begin{aligned}\lambda(\overline{\mathcal{P}}_{\rho'}) \cap \mathcal{Q} &\overset{\mathsf{red}}{\subseteq} \lambda(\mathcal{P}_{x_1})^* \cdot ((\lambda(\mathcal{P}_{x_1,\theta,x_2}) \cdot (\lambda(\overline{\mathcal{P}}_{\rho''}) \cap \mathcal{Q}) \cdot \lambda(\mathcal{P}_{y_1,\chi,y_2})) \cap \mathcal{Q}) \cdot \lambda(\mathcal{P}_{y_2})^* \\ &\overset{\mathsf{red}}{\subseteq} \varpi^* \cdot ((\lambda(\mathcal{P}_{x_1,\theta,x_2}) \cdot \varpi \cdot \lambda(\mathcal{P}_{y_1,\chi,y_2})) \cap \mathcal{Q}) \cdot \varpi^* \\ &\overset{\mathsf{red}}{\subseteq} \varpi^* \cdot \emptyset \cdot \varpi^* = \emptyset,\end{aligned}$$

and therefore that $\lambda(\overline{\mathcal{P}}_\rho) \cap \mathcal{Q} = \emptyset$ as well.

Finally, if $\lambda(\rho)$ is not a prefix of a Dyck word, then the relation $\lambda(\overline{\mathcal{P}}_\rho) \cap \mathcal{Q}_{\mathsf{init}} = \emptyset$ is immediate when ρ itself is not an approximate Dyck path (since $\mathcal{Q}_{\mathsf{init}} \subseteq \mathcal{Q}$). Hence, we assume that ρ is an approximate Dyck path. Let ρ_1 be the longest prefix of ρ such that $\lambda(\rho_1)$ is a prefix of a Dyck word. Let us write ρ as a path of the form $\rho = \rho_1 \cdot (x, \theta, y) \cdot \rho_2$, where (x, θ, y) is an edge of G.

Let Λ be the reduction of $\lambda(\rho_1)$. If $\theta \in \{\ell_1, \ell_2\}$, then the word $\lambda(\rho_1) \cdot \theta \cdot \overline{\theta}$ reduces to Λ, hence is a prefix of a Dyck word, contradicting the maximality of ρ_1. Hence, we know that $\theta \in \{\overline{\ell}_1, \overline{\ell}_2\}$.

Now, let us assume that Λ is not the empty word. Since $\lambda(\rho_1)$ is a prefix of a Dyck word, Λ must contain only letters ℓ_1 and ℓ_2. Hence, without loss of generality, we assume that the rightmost letter of Λ is ℓ_1. If $\theta = \overline{\ell}_1$, then the word $\lambda(\rho_1) \cdot \theta \cdot \ell_1$ reduces to Λ too, which is impossible. Finally, if $\theta = \overline{\ell}_2$, then $\lambda(\rho_1) \cdot \theta$ reduces to $\Lambda \cdot \overline{\ell}_2$, which is not an approximate Dyck word, contradicting the definition of ρ.

Hence, Λ is the empty word, i.e. ρ_1 is a Dyck path. It follows that

$$\begin{aligned}\lambda(\overline{\mathcal{P}}_{\rho_1 \cdot (x,\theta,y)}) \cap \mathcal{Q}_{\mathsf{init}} &\overset{\mathsf{red}}{\subseteq} (((\lambda(\overline{\mathcal{P}}_{\rho_1}) \cap \mathcal{Q}) \cdot \lambda(\mathcal{P}_{x,\theta,y})) \cap \mathcal{Q}_{\mathsf{init}}) \cdot \lambda(\mathcal{P}_y)^* \\ &\overset{\mathsf{red}}{\subseteq} ((\varpi \cdot \lambda(\mathcal{P}_{x,\theta,y})) \cap \mathcal{Q}_{\mathsf{init}}) \cdot \varpi^* \\ &\overset{\mathsf{red}}{\subseteq} \emptyset \cdot \varpi^* = \emptyset,\end{aligned}$$

and therefore that $\lambda(\overline{\mathcal{P}}_\rho) \cap \mathcal{Q}_{\mathsf{init}} = \emptyset$ as well. □

Proposition 16 paves the way to proving Proposition 8. Indeed, if ρ is the nominal ancestor of a Dyck path π in $\mathcal{G}$, then $\lambda(\rho)$ itself is a prefix of a Dyck word (since $\lambda(\pi) \in \lambda(\overline{\mathcal{P}}_\rho) \cap \mathcal{Q}_{\mathsf{init}}$). It remains to prove that $\lambda(\rho)$ must be an actual Dyck word.

The following lemma establishes that the number of occurrences of elements in $\{\ell_1, \ell_2\}$ within the word $\lambda(\rho)$ coincides with the number of occurrences of elements in $\{\overline{\ell}_1, \overline{\ell}_2\}$. This is not sufficient to be a Dyck word, but this will help conclude in Proposition 18.

Lemma 17. *Let π be a Dyck path from s to t in $\mathcal{G}$, let ρ be the nominal ancestor of π, and let $\lambda(\rho) \in L^*$ be the label of ρ. In addition, let $\mu : L^* \mapsto \mathbb{Z}$ be the morphism of monoids defined by $\mu(\ell_1) = \mu(\ell_2) = 1$ and $\mu(\overline{\ell}_1) = \mu(\overline{\ell}_2) = -1$. Then, we have $\mu(\lambda(\rho)) = 0$.*

Proof. Let us identify the monoid $\mathcal{L}^*$ with a sub-monoid of the free group $\mathbb{Z} * \mathbb{Z}$, where the letters 0, $\overline{0}$, 1 and $\overline{1}$ are identified with the elements $(1, 0)$, $(-1, 0)$, $(0, 1)$ and $(0, -1)$ respectively. Then, let Θ be the canonical projection of the free group $\mathbb{Z} * \mathbb{Z}$ onto the free group $\mathbb{Z}_2 * \mathbb{Z}_2$. Let γ be the element $(0, 1) \cdot (1, 0)$ of $\mathbb{Z}_2 * \mathbb{Z}_2$. Observe that γ is not a torsion element of $\mathbb{Z}_2 * \mathbb{Z}_2$, and therefore that the group $\{\gamma^i \mid i \in \mathbb{Z}\}$ is isomorphic with $\mathbb{Z}$.

It comes at once that Θ maps all of the languages ω_+, ω_-, ω, ϖ_+, ϖ_- and ϖ to the singleton set $\{\gamma^0\}$. It follows that $\Theta(\lambda(\mathcal{P}_x)) = \{\gamma^0\}$, that $\Theta(\lambda(\mathcal{P}_{x,\ell_1,y})) = \Theta(\lambda(\mathcal{P}_{x,\ell_2,y})) = \{\gamma\}$ and that $\Theta(\lambda(\mathcal{P}_{x,\overline{\ell}_1,y})) = \Theta(\lambda(\mathcal{P}_{x,\overline{\ell}_2,y})) = \{\gamma^{-1}\}$ for all $x, y \in V$.

Hence, let $v_0, \ldots, v_k$, $\pi_1, \ldots, \pi_k$ and $f_\pi : \{1, \ldots, k\} \mapsto L$ be the nominal vertex sequence, the nominal decomposition and the nominal label mapping of π. Observe that, as an element of $\mathbb{Z} * \mathbb{Z}$, $\lambda(\pi) = \lambda(\pi_1) \cdot \ldots \cdot \lambda(\pi_k)$ is the neutral element of the group, whence $\gamma^0 = \Theta(\lambda(\pi_1)) \cdot \ldots \cdot \Theta(\lambda(\pi_k))$.

We showed above that $\Theta(\lambda(\pi_i)) = \gamma^0$ for all $i \in \{1, \ldots, k\} \setminus \mathsf{dom}(f_\pi)$ and that $\Theta(\lambda(\pi_i)) = \gamma^{\mu(f_\pi(i))}$ for all $i \in \mathsf{dom}(f_\pi)$, which completes the proof. □

The above lemmas lead to the following result, which, together with Proposition 9, completes the proof of Proposition 8.

Proposition 18. *Let π be a Dyck path from s to t in $\mathcal{G}$. There exists a Dyck path from s to t in G.*

Proof. Let ρ be the nominal ancestor of π. Let $\lambda(\rho)$ be the label of ρ and let Λ be the reduction of $\lambda(\rho)$. Lemma 17 proves that $\mu(\lambda(\rho)) = \mu(\Lambda) = 0$. Proposition 16 proves that $\lambda(\rho)$ is a prefix of a Dyck word. It follows that Λ has only letters ℓ_1 and ℓ_2, and therefore that $|\Lambda| = \mu(\Lambda) = 0$. This means that Λ is the empty word, i.e. that $\lambda(\rho)$ is a Dyck word, which completes the proof. □